

**APPLICATION
FOR
UNITED STATES LETTERS PATENT**

APPLICANT NAME: Durham et al.

TITLE: Computerized System, Method and Program Product for
Managing an Enterprise Storage System

DOCKET NO.: END920030129US1

INTERNATIONAL BUSINESS MACHINES CORPORATION

CERTIFICATE OF MAILING UNDER 37 CFR 1.10

I hereby certify that, on the date shown below, this correspondence is being deposited with the United States Postal Service in an envelope addressed to Mail Stop Patent Application, Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450 as "Express Mail Post Office to Addressee" Mailing Label No. EL102609375US

on 11/20/2003

Georgia Y. Brundage

Name of person mailing paper

Georgia Y. Brundage
Signature

11.20.03
Date

COMPUTERIZED SYSTEM, METHOD AND PROGRAM PRODUCT FOR MANAGING AN ENTERPRISE STORAGE SYSTEM

FIELD OF THE INVENTION

[0001] The present invention generally relates to a computerized system, method and program product for managing an enterprise storage system. Specifically, the present invention optimizes the use the storage cells of an enterprise storage system.

BACKGROUND OF THE INVENTION

[0002] As the use of information technology (IT) becomes more prevalent, it has become increasingly common for many organizations to use enterprise storage systems for storing files and the like. In an enterprise storage system, multiple data storage cells or farms are provided at different geographic locations throughout the country or world. These cells can be used for sharing files among individuals within the organizations. For example, International Business Machines Corp. of Armonk, NY uses a technology known as the Global Storage Architecture (GSA) for file sharing.

[0003] Unfortunately, in many enterprise storage systems such as the GSA, vast quantities of the available resources are unbalanced and/or unused due to the clustering design. Specifically, it is currently the case that files are stored based on geographic parameters. For example, if user "A" is located on the east coast of the United States, currently technologies would store his/her files in the nearest cell. However, such storage might not make the best use of the overall system. Rather, it could be the case that the

east coast cell is nearing its capacity, while a midwest cell is virtually empty. Further, it could be the case that the east coast cell does not support a needed rate of network connectivity, while another cell does. Accordingly, none of the existing technologies take performance-based parameters such as storage capacity, user requirements, etc., into consideration when selecting a file sharing location.

[0004] In view of the foregoing, there exists the need for a computerized system, method and program product for managing enterprise storage farms. Specifically, a need exists for a system that better utilizes storage resources. To this extent, a need exists for a system that stores files based on certain performance parameters.

SUMMARY OF THE INVENTION

[0005] In general, the present invention provides a computerized system, method and program product for managing an enterprise storage system. Specifically, under the present invention a storage grid manager receives requests for storing files from users. Based on certain predetermined performance parameters, the storage grid manager identifies appropriate/selected storage cells of the enterprise storage system for storing the files. Thereafter, the storage grid manager routes the requests to storage grid controllers associated with the selected storage cells. The storage grid controllers control access to the storage cells to ensure that only authorized access is granted. Once the files are stored, the storage grid manager will update a mapping to identify the precise storage cells in which each file is stored. Therefore, when a user issues a request to retrieve a

file, the storage grid manager consults the mapping, identifies the corresponding storage cell, and retrieves the file.

[0006] A first aspect of the present invention provides a computerized system for managing an enterprise storage system, comprising a storage grid manager for receiving requests for storing files from users, and for routing the requests to storage grid controllers associated with selected storage cells of the enterprise storage system where the files will be stored, wherein the selected storage cells are identified based on at least one predetermined performance parameter.

[0007] A second aspect of the present invention provides a computerized system for managing an enterprise storage system, comprising: a storage grid manager for receiving requests for storing files from users; and a set of storage grid controllers associated with a set of storage cells of the enterprise storage system, wherein the storage grid manager identifies selected storage cells for storing the files based on at least one predetermined performance parameter and routes the requests to the storage grid controllers associated with the selected storage cells.

[0008] A third aspect of the present invention provides a computer-implemented method for storing files in an enterprise storage system, comprising: receiving requests on a storage grid manager to store the files; identifying storage cells of the enterprise storage system for storing the files based on at least one performance parameter; routing the requests from the storage grid manager to storage grid controllers associated with the storage cells; and storing the files in the storage cells.

[0009] A fourth aspect of the present invention provides a computer-implemented method for retrieving files from an enterprise storage system, comprising: receiving requests on a storage grid manager to retrieve the files; identifying storage cells of the enterprise storage system in which the files are stored based a mapping; routing the requests from the storage grid manager to storage grid controllers associated with the storage cells; and retrieving the files from the storage cells.

[0010] A fifth aspect of the present invention provides a program product stored on a recordable medium for managing an enterprise storage system, which when executed, comprises a storage grid manager for receiving requests for storing files from users, and for routing the requests to storage grid controllers associated with selected storage cells of the enterprise storage system where the files will be stored, wherein the selected storage cells are identified based on at least one predetermined performance parameter.

[0011] Therefore, the present invention provides a computerized system, method and program product for managing an enterprise storage system.

BRIEF DESCRIPTION OF THE DRAWINGS

[0012] These and other features of this invention will be more readily understood from the following detailed description of the various aspects of the invention taken in conjunction with the accompanying drawings in which:

[0013] Fig. 1 depicts a storage grid architecture according to the present invention.

[0014] Fig. 2 depicts the relationship between the storage grid manager, the storage grid controllers and the storage cells of Fig. 1 in greater detail.

[0015] Fig. 3 depicts an access control specification according to the present invention.

[0016] Fig. 4 depicts a more detailed diagram of a storage grid manager of Fig. 1.

[0017] Fig. 5 depicts a state diagram for storing files according to the present invention.

[0018] Fig. 6 depicts a state diagram for retrieving files according to the present invention.

[0019] The drawings are merely schematic representations, not intended to portray specific parameters of the invention. The drawings are intended to depict only typical embodiments of the invention, and therefore should not be considered as limiting the scope of the invention. In the drawings, like numbering represents like elements.

BEST MODE FOR CARRYING OUT THE INVENTION

[0020] As indicated above, the present invention provides a computerized system, method and program product for managing an enterprise storage system. Specifically, under the present invention a storage grid manager receives requests for storing files from users. Based on certain predetermined performance parameters, the storage grid manager identifies appropriate/selected storage cells of the enterprise storage system for storing the files. Thereafter, the storage grid manager routes the requests to storage grid controllers associated with the selected storage cells. The storage grid controllers control access to the storage cells to ensure that only authorized access is granted. Once the files are stored, the storage grid manager will update a mapping to identify the precise storage cells in which each file is stored. Therefore, when a user issues a request to retrieve a

file, the storage grid manager consults the mapping, identifies the corresponding storage cell, and retrieves the file.

[0021] It should be understood in advance that although the invention is typically used to manage the storage and retrieval “files,” it can be used to manage the storage and retrieval of any type of information.

[0022] In any event, referring now to Fig. 1 a storage grid architectural (SGA) diagram 10 according to the present invention is depicted. In general, SGA 10 allows users such as user 12 to store files and the like in one or more storage cells/farms 14A-D of enterprise storage system 16. To this extent, SGA is typically implemented over a network such as the Internet, a local area network (LAN), a wide area network (WAN), a virtual private network (VPN), etc. As such, communication between the various components shown in Fig. 1 could occur in a client-server or server-server environment via a direct hardwired connection (e.g., serial port), or via an addressable connection that may utilize any combination of wireline and/or wireless transmission methods. The components may utilize conventional network connectivity, such as Token Ring, Ethernet, WiFi or other conventional communications standards. Moreover, connectivity could be provided by conventional TCP/IP sockets-based protocol. In this instance, the components could utilize an Internet service provider to establish connectivity. It should be appreciated that the architecture shown in Fig. 1 is intended to be illustrative only and that other variations could be implemented within the scope of the present invention. For example, any quantity of storage cells 14A-D and/or storage grid managers 24 could be provided.

[0023] Regardless, in a typical embodiment, user 12 will store a file by generating a request via an application on his/her computer system (e.g., a personal computer, a workstation, a laptop, a hand-held device, etc.). Depending on the desired communication protocol, the request will be transmitted to web server 18 and/or application server 20, and from there to a respective "protocol" server 22A-C. For example, if user 12 desired to use File Transfer Protocol (FTP) to request storage of the file, the request would be routed to web server 18 and then to FTP server 22A. As such, the present invention can accommodate any type of protocol user 12 might wish to utilize. For example, storage can be provided for Hypertext Transfer Protocol (HTTP), Common Internet File System (CIFS), Network File System (NFS), Server Message Block Protocol (SMB), etc.

[0024] Once the request is received by a protocol server 22A-C, it will be routed to storage grid manager 24 (e.g., to one a possible group of storage grid managers). As will be further described below, storage grid manager 24 is configured to determine the most optimal or appropriate storage cell 14A-D for storing the file. In general, this determination is made under the present invention based on certain predetermined performance parameters. For example, the determination can be made based upon a user identity, a storage cost, a specific user requirement, desired a cell usage pattern/balance, a security requirement, a storage cell availability, a redundancy requirement and a network optimization requirement. Thus, if a file required access by a group of users, it could be stored in a storage cell 14A-D that is most central to all such users, not just the user requesting the storage. Moreover, if a certain storage cell 14A-D was nearing capacity,

the file might be stored in another (albeit more distant) storage cell 14A-D so that the load of storage cells 14A-D could be kept balanced.

[0025] In any event, once the appropriate storage cell 14A-D is identified, storage grid manager 24 will route the request to the storage grid controller 26A-D corresponding thereto. Specifically, each storage cell 14A-D is associated with one storage grid controller 26A-D. This relationship is further depicted in Fig. 2. As can be seen each storage grid controller 26A-D is associated with a single storage cell 14A-D. This relationship allows each storage grid controller 14A-D to serve at least two primary functions. Namely, to provide resource information about storage cells 14A-D to storage grid manager 24, and to act as a “gatekeeper” of its respective storage cell 14A-D. In the case of the former, under the present invention, storage grid manager 24 will receive information regarding capacity of the storage cells 14A-D. This information can be communicated upon request by storage grid manager 24, or updates can be periodically sent from each storage cell 14A-D via storage grid controllers 26A-D. In either event, the resource information is used by storage grid manager 24 in determining which storage cell 14 is the most appropriate for storing the file. For example, as indicated above, if storage cell 14A is closest geographically to user 12, but is close to capacity, storage grid manager 24 might route the request to storage cell 14B.

[0026] Assuming storage cell 14B was identified as the most appropriate storage location for the file, storage grid manager 24 will route the request to store the file to storage grid cell 14B via storage grid controller 26B. In so doing, storage grid manager 24 will likely add additional information such as an identification and/or authorization/authentication

information for user 12. Moreover, the request will typically be routed from storage grid manager 24 using the same protocol in which it was received by storage grid manager 24, or via a private protocol.

[0027] Upon receipt of the request, the storage grid controller 26B will consult access manager 28, which will access PD Lightweight Directory Access Protocol (LDAP) directory 30 to determine if user 30 is authorized to access storage cell 14B (i.e., enforce access control specifications for storage cells 14A-D). In a typical embodiment, LDAP directory 30 is a DB2 database that stores permissions in an additive hierarchy such as shown in Fig. 3. Specifically, after a user is given permission to a path, the user has that same permission to all sub-paths under it. The user cannot be added with the same permission to any of the sub-paths. Ones and zeros represent permissions in the database. A one indicates that the user has a particular permission to a resource and a zero is a place holder. A zero does not indicate a negative permission. Rather, it merely says that a particular permission was not granted at that level of the hierarchy to a particular user. Moreover, there can be a situation where a user has a record granting access at a certain level and a level below that. This situation can occur when access to the higher level was granted after access to the lower level had been granted.

[0028] Assuming that user 12 was authorized to access storage cell 14B, the file will be stored therein and storage grid manager 24 will update a mapping that correlates files with the storage cells in which they are stored. Specifically, among its other functions, storage grid manager 24 maintains a mapping that it updates as files are stored in storage cells 14A-D. Thus, if user 12 later issues a request to retrieve the file, the request will be

routed to storage grid manager 24, which will consult the mapping to identify the appropriate storage cell 14B. Upon so doing, storage grid manager 24 will route the request to storage cell 14B and/or storage grid controller 26B, which will again perform access control via access manager 28 and LDAP 30. Assuming user 12 is still authorized to access storage cell 14B, the file will be retrieved for access by user 12.

[0029] Referring now to Fig. 4, a more detailed diagram of storage grid manager 24 is shown. As depicted, storage grid manager 24 generally includes central processing unit (CPU) 40, memory 42, bus 44, input/output (I/O) interfaces 46, external devices/resources 48 and storage unit 49. CPU 40 may comprise a single processing unit, or be distributed across one or more processing units in one or more locations, e.g., on a client and computer system. Memory 42 may comprise any known type of data storage and/or transmission media, including magnetic media, optical media, random access memory (RAM), read-only memory (ROM), a data cache, etc. Moreover, similar to CPU 40, memory 42 may reside at a single physical location, comprising one or more types of data storage, or be distributed across a plurality of physical systems in various forms.

[0030] I/O interfaces 46 may comprise any system for exchanging information to/from an external source. External devices/resources 48 may comprise any known type of external device, including speakers, a CRT, LCD screen, handheld device, keyboard, mouse, voice recognition system, speech output system, printer, monitor/display, facsimile, pager, etc. Bus 44 provides a communication link between each of the components in storage grid manager 24 and likewise may comprise any known type of transmission link, including electrical, optical, wireless, etc.

[0031] Storage unit 49 can be any system (e.g., database) capable of providing storage for information under the present invention. Such information could include, for example, the mapping maintained by storage grid manager 24. As such, storage unit 49 could include one or more storage devices, such as a magnetic disk drive or an optical disk drive. In another embodiment, storage unit 49 includes data distributed across, for example, a local area network (LAN), wide area network (WAN) or a storage area network (SAN) (not shown). Furthermore, although not shown, additional components, such as cache memory, communication systems, system software, etc., may be incorporated into storage grid manager 24. In addition, it should be appreciated that although not shown, the computer system operated by user, the web server, the application server, the protocol servers, the access manager and the storage grid controllers 26A-D can include computerized components similar to storage grid manager 24. Such components have not been shown for brevity purposes.

[0032] In any event, shown in memory 42 of storage grid manager 24 as a program product is storage management system 50, which generally includes cell identification system 52, resource information system 54, request routing system 56, mapping system 58 and notification system 60. Under the present invention, requests to store files will be received by storage grid manager 24. At that point, cell identification system 52 will use the performance parameters to identify appropriate storage cells 14A-D for storing the files. As indicated above, such parameters could include, among others, a user identity, a storage cost, a user requirement, a desired cell usage pattern/balance, a security requirement, a storage cell availability, a redundancy requirement and a network

optimization requirement. In making this determination, storage grid manager 24 will use resource information about storage cells 14A-D such as available space that is received via resource information system 54. This information can be obtained based on queries made by resource information system 54, or via updates periodically sent via storage grid controllers 26A-D.

[0033] In any event, once the appropriate storage cells 14A-D are identified, request routing system 56 will route the requests to the appropriate storage cells 14A-D (i.e., via the storage grid controllers 26A-D associated therewith). Upon receipt, access enforcement systems (not shown) within the storage grid controllers 26A-D will determine whether the users are authorized to access the respective storage cells 14A-D. If so, the files are stored accordingly. Upon such storage, mapping system 58 will update the mapping to reflect the locations in which the files were stored, and notification system 60 will notify the users accordingly.

[0034] Requests to retrieve the files will similarly be received by storage grid manager 24. Upon receipt, cell identification system 52 will consult the mapping to identify the pertinent storage cells 14A-D, and request routing system 56 will route the retrieval requests to the identified storage cells 14A-D. Similar to storage of the files, access control systems within the storage grid controllers 26A-D will ensure that the users are authorized to have such access. If so, the files are retrieved and notification system 60 will notify the users accordingly.

[0035] It should be understood that the present invention can be realized in hardware, software, or a combination of hardware and software. Any kind of computer system(s) -

or other apparatus adapted for carrying out the methods described herein - is suited. A typical combination of hardware and software could be a general purpose computer system with a computer program that, when loaded and executed, carries out the respective methods described herein. Alternatively, a specific use computer, containing specialized hardware for carrying out one or more of the functional tasks of the invention, could be utilized. The present invention can also be embedded in a computer program product, which comprises all the respective features enabling the implementation of the methods described herein, and which - when loaded in a computer system - is able to carry out these methods. Computer program, software program, program, or software, in the present context mean any expression, in any language, code or notation, of a set of instructions intended to cause a system having an information processing capability to perform a particular function either directly or after either or both of the following: (a) conversion to another language, code or notation; and/or (b) reproduction in a different material form.

[0036] Referring now to Fig. 5, a state diagram 100 of file storage is shown. As depicted, when a user wishes to store a file, the user will first submit a store request to the storage grid manager. Using resource information obtained from the storage cells (via the storage grid controllers), the storage grid manager will identify the most appropriate storage cell in which to store the file. Under the present invention, this determination is based on factors in addition or other than geography. For example, a user identity, a storage cost, a user requirement, a desired cell usage pattern/balance, a security requirement, a storage cell availability, a redundancy requirement, a network optimization requirement, etc.

Once identified, the storage grid manager will forward the request to the associated storage grid controller, which after verifying authorization of the user, will permit storage of the file. Upon completion, the storage grid manager will update the mapping and notify the user.

[0037] Referring to Fig. 6 a state diagram 200 of file retrieval is shown. As depicted, when a user wishes to retrieve the file, the user will first submit a retrieval request to the storage grid manager. Using the mapping, the storage grid manager will identify the appropriate storage cell where the requested file is stored, and communicate the request thereto (via the associated storage grid controller). After confirming authorization of the user to access the file/storage cell, the file will be retrieved for the user and a corresponding notification will be sent.

[0038] The foregoing description of the preferred embodiments of this invention has been presented for purposes of illustration and description. It is not intended to be exhaustive or to limit the invention to the precise form disclosed, and obviously, many modifications and variations are possible. Such modifications and variations that may be apparent to a person skilled in the art are intended to be included within the scope of this invention as defined by the accompanying claims.